

REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-00-

98

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Project, Washington, DC 20503.

0030

ated, and completing an
Services, Directorate fo
503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
		01 Jun 96 to 30 Nov 98 Final
4. TITLE AND SUBTITLE The application of free-space optical interconnects to high performance communication decoding		5. FUNDING NUMBERS 62173C 1651/01
6. AUTHOR(S) Dr Haney		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) George Mason University 4400 University Drive Fairfax VA 22030-444		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NE 801 North Randolph Street Rm 732 Arlington, VA 22203-1977		10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-96-1-0282
11. SUPPLEMENTARY NOTES		
12a. DISTRIBUTION AVAILABILITY STATEMENT APPROVAL FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED		12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words) Development of the "Two-bounce" interconnection concept, that is based on topological transformations, and which minimizes the smart pixel resource required for the Viterbi architecture and has wider implications for free-space optical interconnects. Completion of the initial optomechanical evaluation system for the Viterbi system which used fiber-coupled arrays to simulate the eventual smart pixel array I/O. Development of general scaling laws for free-space optical systems, which show the size, volume, power consumption and latency benefits of free-space optics for high bisection bandwidth interconnection applications, such as the Viterbi application. Development of a novel hybrid macro/micro-optical scheme that simplifies the optical design while minimizing critical aberrations in the system. The following sections detail the key results for each of these progress area. The progress made in this program has resulted in 3 refereed journal papers and 6 conference papers.		
		16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED
		20. LIMITATION OF ABSTRACT UL

The Application of Free-space Optical Interconnects to High Performance Communication Decoding

Grant No. F49620-96-1-0282

Final Technical Report

Covering the period: June 1, 1996 – November 30, 1998

Submitted to: AFOSR/NE

20000214 037

Principal Investigator: Dr. Michael W. Haney
George Mason University, ECE Dept.
Mail stop 1G5, Fairfax, VA, 22030-4444
Phone: (703)993-157, Fax: (703)993-1601
E-mail: mhaney@gmu.edu

Table of Contents

I. Executive Summary	1
II. "Two-bounce" Interconnection Concept	2
A. Background	2
B. Two-Bounce Architecture	3
C. Interconnection Pattern Examples	7
III. Optomechanical Evaluation System	9
IV. Scaling Laws for Free-space Optical Interconnection Systems	11
A. Motivation	11
B. Electrical Interconnection Requirements	12
C. Optical Interconnection Requirements	15
D. Discussion	18
V. Hybrid Macro/Micro-optical Interconnection Concept	22
VI. Conclusions	28
VII References	29

I. Executive Summary

This report covers the theoretical and experimental advancements made during the 2.5 years research effort on the application of free-space optical interconnection techniques to high performance communications decoding. VLSI implementations of the elegant and powerful Viterbi convolutional decoding algorithm (VA), which uses a recursive parallel search computation, are limited by the massive intra- and inter-chip communications requirements between nodes of the search graph. This constraint limits the number of states (nodes of the VA graph), for high-speed applications, and hence the overall performance of the VA. The approach being implemented in this research program overcomes the VA 2-D communications bottleneck by combining the rapidly emerging smart pixel technology with 3-D folded free-space optical interconnects (FSOI) to implement the required interconnection network. The interconnection densities provided by FSOI and smart pixel technology provide the potential for orders of magnitude improvements in bit error rates (BER) and speed with an order of magnitude reduction in size/weight/power requirements for high performance receivers. This concept thus significantly expands the application domain of the powerful VA to platforms that otherwise could not support processors based on larger and power hungry conventional metallic interconnection technology. The concept developed in this program leverages free-space optical and smart-pixel technologies that are being developed for telecommunication switches and parallel computer networks. This report highlights the significant progress made in four areas:

- ◆ Development of the "Two-bounce" interconnection concept, that is based on topological transformations, and which minimizes the smart pixel resource required for the Viterbi architecture and has wider implications for free-space optical interconnects.
- ◆ Completion of the initial optomechanical evaluation system for the Viterbi system which used fiber-coupled arrays to simulate the eventual smart pixel array I/O.
- ◆ Development of general scaling laws for free-space optical systems, which show the size, volume, power consumption and latency benefits of free-space optics for high bisection bandwidth interconnection applications, such as the Viterbi application.
- ◆ Development of a novel hybrid macro/micro-optical scheme that simplifies the optical design while minimizing critical aberrations in the system..

The following sections detail the key results for each of these progress areas. The progress made in this program has resulted in 3 refereed journal papers and 6 conference papers.

II. Two-Bounce Interconnection Concept

A. Background

Free-space optical interconnections (FSOI) have been shown to overcome communications limitations in large, globally interconnected multi-processor architectures by scaling well for the multi-Terabit bisection bandwidth regime [1,2]. Several *macro*-optical approaches to shuffle interconnection networks have been proposed and demonstrated [3-10]. There appears, however, to be a significant trade-off between the fundamental scaling benefits of 3-D free-space macro-optical approaches and the inherent arbitrary interconnection flexibility of space variant *micro*-optical interconnection approaches. While multi-chip macro optical interconnection approaches, such as the one shown in Figure 1, have been shown to scale effectively to high bisection bandwidth problems, they are limited, by their high degree of space-invariance, to implementing only regular shuffle link patterns. A macro-optical interconnection approach is desired which provides *arbitrary* interconnections, yet retains the beneficial scaling properties of macro-optics.

It is commonly assumed that using higher order k -shuffle based optical interconnections will require the use of $k \times k$ crossbar switches for the local switching elements to achieve arbitrary link patterns. However, as shown in this paper, 3-D topological transformations make it possible to avoid the use of $k \times k$ crossbar switches entirely, while requiring only the minimum number of 2×2 switching elements. The Two-Bounce architecture achieves, without changing lens positions or attributes, a completely arbitrary interconnection pattern – through changing only local 2×2 switch electronic interconnections. Furthermore, the optical system can be implemented with a symmetric macro-optical multi-chip arrangement, thereby allowing the interconnection to be folded back onto itself in a reflective single plane architecture that achieves the required high degree of opto-mechanical alignment [10].

The application of FSOI techniques to a multi-processor interconnection problem can be viewed as a mapping of the network's functional interconnection pattern onto a 3-D optical interconnection architecture [13]. Such a mapping amounts to a topological transformation, which preserves the interconnection pattern, and functionality of the architecture's configuration, but achieves performance advantages owing to the use of 3-D space and smart pixel capabilities. In fact the architecture can be represented as a series of topological transformations that each exploit a performance advantage of photonic interconnects. The cumulative performance advantage of a FSOI implementation of a network architecture, therefore, derives from the aggregate advantages of several distinct topological transformations of the link interconnection pattern.

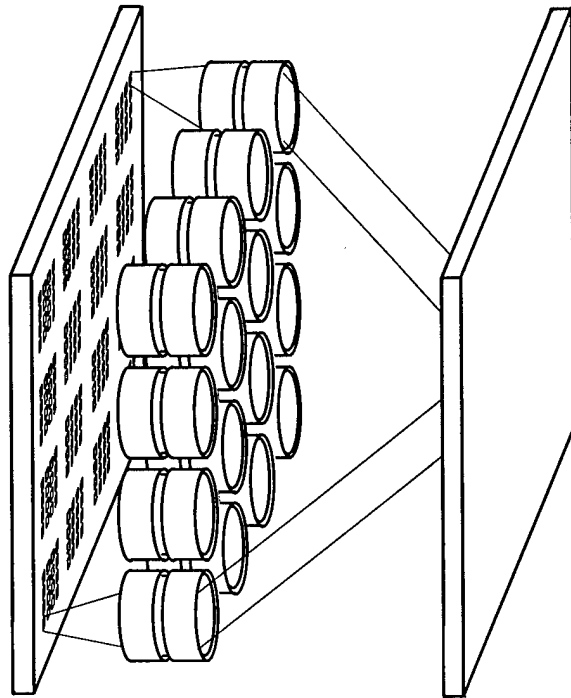


Figure 1. Schematic depiction of reflective macro-optical multi-chip interconnection module.

Examples of topological transformations that apply to FSOI banyan-based networks and the motivation for using them are shown in Figure 2. Figure 2a depicts the isomorphism between banyans consisting of butterflies and shuffles. Using an optical shuffle link pattern between stages of the banyan simplifies the optical design and facilitates further transformations as described below. Figure 2b shows the formatting of the shuffle as a 2-D shuffle, rather than a 1-D shuffle – to take better advantage of optical and MCM packaging techniques. Arraying the smart pixel on self-similar grids (Figure 2c), rather than rectilinear grids increased the multi-chip pixel density and optical efficiency [14]. Figure 2d depicts the spatial interleaving of multiple stages – to cluster nodes and thereby reduce the amount of required electronic resources in the smart pixel [15,16]. Furthermore, if every stage is a shuffle, then this topological transformation enables the use of a single reflective optical system. Figure 2e shows this common plane reflective approach – to distribute the smart pixels across a single backplane, simplify optical alignment, and reduce the number of output drivers required [10]. Each of these FSOI topological transformations is motivated by a packaging advantage that leads to a performance enhancement or packaging simplification. The performance enhancements achieved by these topological transformations are made practical only through the use of 3-D FSOI.

B. Two-Bounce Architecture

The previous section described the topological transformations which map regular shuffle interconnected multistage interconnection networks (MINs) onto optical interconnection modules like that of Figure 1. The Benes network is a regular modulo-2 MIN-based network that achieves

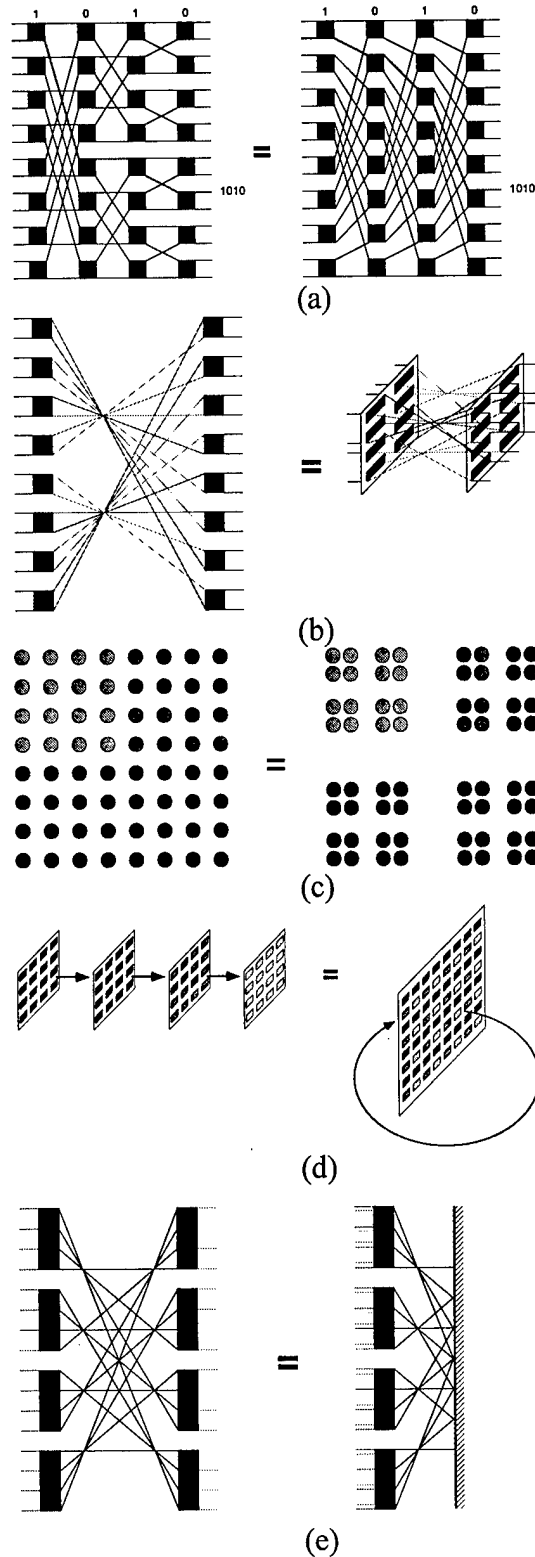


Figure 2. Example topological transformations of multi-stage FSOI architectures.

arbitrary rearrangeable non-blocking interconnections with the minimum number of switching resources [17]. But can the Benes network be implemented with higher order k -shuffle optical modules, without paying the increased switching penalty associated with higher order $k \times k$ crossbars? As discussed below, the Two-Bounce architecture achieves exactly this result through the judicious application of topological transformations that can be implemented with the reflective k -shuffle FSOI module. The topological transformations rearrange the interconnections required for the Benes network resulting in 2 stages of global interconnections, performed optically, and multiple stages of local electronic interconnections. While the Two-Bounce architecture retains the Benes' minimum number of switching resources for an arbitrary permutation network, it also minimizes the global interconnection requirements thereby minimizing the FSOI interconnect resource requirement. The Benes network and the topological transformations applied to it are discussed in sections 3A to 3B below.

Benes Architecture

The Benes network, shown in Figure 3, consists of back to back butterfly networks. The resulting network consists of $2\log_2(N)-1$ switching stages and $2\log_2(N)-2$ interconnection stages, where N is the number of nodes. As depicted in Figure 3, the first butterfly interconnection is oriented in a forwards direction, whereas the second butterfly interconnection is oriented in the reverse direction. This network has been shown to require the minimum number of 2×2 switching elements to effect a rearrangeable non-blocking permutation network [17]— any permutation of inputs to outputs can be realized with this relatively simple switching network. The simplicity of the network has its price — the routing algorithm for the Benes network requires global information of the permutation and is iterative, and therefore does not readily lend itself directly to low latency packet switching applications. However, the Benes network is useful for networks that can use out-of-band reconfiguration or which can store a precompiled set of interconnection patterns. For example, FFTs used in digital signal processing can be implemented on multiprocessor architectures in which the processors are linked by the butterfly patterns required by the FFT algorithm. 2-D FFT implementations, used for example in Synthetic Aperture Radar (SAR), also require a memory corner-turn interconnection that amounts to a transpose of the data. These types of interconnections are notoriously difficult due to their high bisection bandwidth. The Two-Bounce architecture is particularly well suited to these types of interconnections because it can

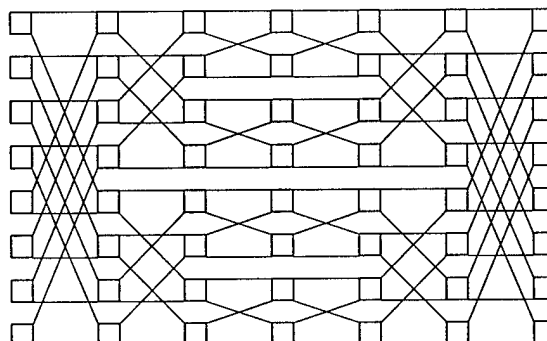


Figure 3. Butterfly based Benes Network for $N=16$.

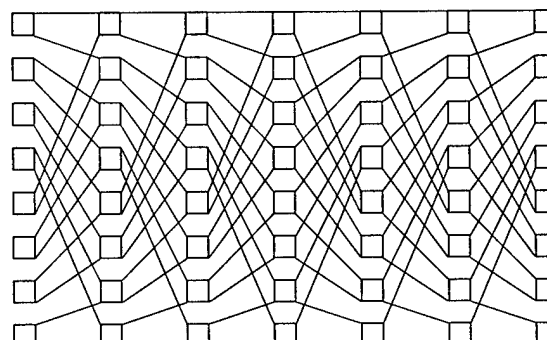


Figure 4. Perfect shuffle based Benes Network for $N=16$ which is isomorphic to the network shown in Figure 3.

pre-store the required switch settings for each stage of the FFT's butterfly, as well as the corner-turn settings.

As described in Section II, the butterfly interconnection network is isomorphic to $\text{Log}_2 N$ shuffle interconnections as shown in Figure 2a. The application of this topological transformation results is a new shuffle based Benes network, depicted in Figure 4, which is comprised of identical shuffle interconnections between switching elements. The identical plane-to-plane interconnection patterns make possible another topological transformation in which the interconnection module is interleaved and folded back onto itself. However, at this point, the Benes network is still comprised of $\sim 2\text{Log}_2 N - 2$ stages of shuffles, each requiring global interconnections resources. The scaling benefits of macro optics are best utilized when the optical interconnection pattern is global between multiple chips, i.e., higher order shuffles corresponding to the number of OEICs interconnected in the architecture [1,2]. This motivates the transformation of the Benes network into an architecture utilizing higher order shuffles.

Topological Transformation of 2-Shuffles into Higher Order Shuffles

A perfect shuffle is a global interconnection pattern that amounts to a 1-bit rotation of an address [11]. A shuffle-exchange stage consists of a perfect shuffle followed by a set of $N/2 \times 2$ exchange bypass switches, where N is the number of nodes. Therefore, a series of M shuffle-exchange stages performs a sequence of M rotations, after each of which, the locally connected bypass-exchange switch causes the least significant bit to remain unchanged or switched to its complement. This network can be topologically transformed into a single global $k=2^M$ shuffle followed by routing and switching among the M least significant bits. This makes sense because an M stage 2-shuffle MIN performs the same function as a $k=2^M$ shuffle based MIN that performs M left rotations (in one step) followed by a single set of N/M banyans of size M to set the M least significant bits [18]. This transformation is depicted in Figure 5. Figure 5a depicts two 2-shuffle stages of 16 nodes, where the switching elements are labeled for reference. Figure 5b depicts a single 4-shuffle on the same 16 nodes, with the resultant node labeling. The transformation from Figures 5a to 5b moved only the switching elements, retaining the original interconnections between them. In this fashion, any M 2-shuffle stages can be transformed into a single 2^M global shuffle followed by local routing and switching (amounting to a banyan).

This transformation of a 2-shuffle based architecture into a higher order shuffle based architecture facilitates the mapping of the Benes network onto a k -shuffle optical module, where k is a power of 2. In fact, $k=N^{1/2}$ is the optimum choice for implementing the network on reflective folded modules [19,20] such as Figure 1 because the resulting k -shuffle is symmetric [9]

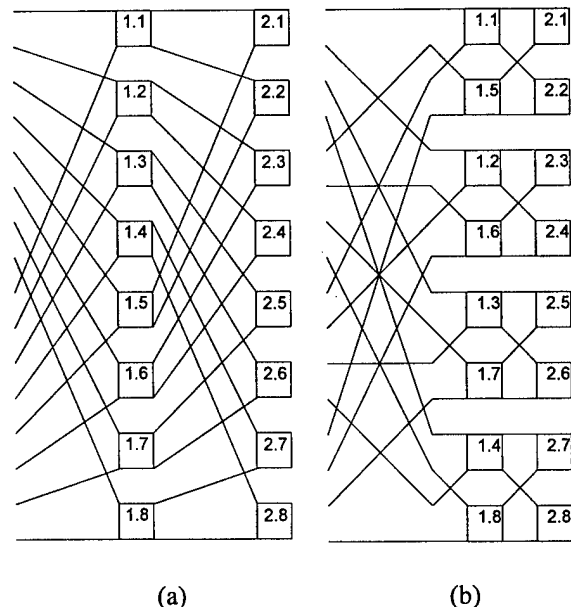


Figure 5. (a) is 2 perfect shuffle exchange stages for $N=16$. (b) is a topological transformation of (a) to a single global 4-shuffle followed by a 4 banyans of 4 elements each.

(i.e., the shuffle rotates half of the bits). Since the equivalence of 2-shuffle mappings to higher order mappings requires an initial shuffle pattern (as shown in Figure 5a), the shuffle based Benes depicted in Figure 4 must be modified to include this initial shuffle pattern to the first and last stages of the Benes. Figure 6 shows the modified 2-shuffle Benes network, with the initial and final shuffles shown as dashed lines. Figure 7 depicts the result of transforming Figure 6 to utilize higher order shuffle interconnections. Note that the resultant architecture also has initial and final k-shuffles ($k=4$ in this example), again shown in dashed lines. Figure 7 is completely equivalent to Figure 6 – it contains the same number of switching elements and they are all interconnected in the same pattern. When a module is built to realize the architecture in Figure 7, the initial and final global interconnections (dashed lines) are not required. The dashed lines only define a mapping between the inputs of Figures 6 and 7, used to determine the switch settings. To implement a mapping of permutation A to permutation B in the interconnection module depicted in Figure 7, the 2-shuffle Benes is solved for mapping of A^* to B^* , where A^* and B^* are defined as follows:

$$A^* = \left[[A]^{-4} \right]^2,$$

$$B^* = \left[[B]^{-4} \right]^2,$$

where the -4 exponent represents an inverse 4-shuffle of the pattern and the 2 exponent represents a 2-shuffle of the pattern. Once the switch settings are determined utilizing A^* and B^* for the 2-shuffle implementation they are directly applied to the higher order implementation and the dashed lines are not needed, and are therefore not implemented.

Even though the global interconnection pattern is implemented with a higher order k-shuffle, the Benes network remains, *logically*, a 2-shuffle Benes implementation. There are still $2\log_2 N - 1$ switching stages and $2\log_2 N - 2$ interconnection stages, only now all but 2 of the interconnection stages are local interconnections. The 2 global interconnections are symmetric optical shuffles with shuffle order (k) equal to $N^{1/2}$. Note that the local electronic routing and switching in the middle switching plane is identical to N/k Benes network, each containing k elements. The first and last electronic switching and routing planes are each comprised of simply N/k k-banyans, because they are not required to perform all permutations within the Benes structure. *The result of mapping the 2-shuffle Benes network onto a higher order shuffle, while retaining the 2×2 switching, results in fewer switching resources than had the Benes network been constructed of higher order shuffles, which would have required higher order $k \times k$ crossbars at each of the 3 switching stages.* Again the resultant architecture is comprised of symmetric shuffles which facilitate the folding of the optical systems and the interleaving of resources into a module such as the one depicted in Figure 1.

C. Interconnection Pattern Examples

In order to illustrate the steps involved for the Two-Bounce arbitrary permutation architecture, two example interconnection patterns with differing requirements are presented. The two interconnection patterns are a matrix transpose, and a Folded Perfect Shuffle [7]. For illustrative purposes the Two-Bounce interconnections is shown in 6 steps: 1) original data positions, 2) data positions after local electronic routing and switching, 3) data after the first global optical interconnection, 4) data after the second stage of local electronic routing and switching, 5) data after the second global optical interconnection and finally, 6) final data positions. To make the example easy to follow, a simple 2×2 chip array is utilized with 2×2 data positions within each chip, corresponding to a data set of 16 nodes. Figure 8 is the Two-Bounce interconnection effecting a transpose of the original data in a matrix fashion. Note that data remains within "chip" boundaries during local routing operations (between stages 1-2, 3-4, and 5-6). The global optical interconnections take place between stages 2-3 and stages 4-5, and are fixed for this, and all, interconnection patterns using the Two-Bounce architecture.

Figure 9 shows a Two-Bounce interconnection effecting the Perfect Shuffle, in this case Folded, of the data set. The global optical interconnection stages of Figure 9 are identical to that of Figure 8. This is a key feature of the Two-Bounce architecture -- the optical interconnection module is fixed. No modification is required to change the interconnection pattern. Only the local electronic routing is changed to modify a transpose interconnection to a Perfect Shuffle interconnection. While the resulting optical interconnection of the Two-Bounce architecture is a Folded Perfect Shuffle, the optical interconnection module is physically different. It contains two lens planes, arranged in a symmetric fashion, facilitating the folding of the

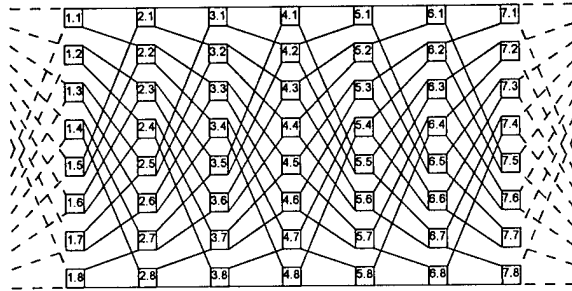


Figure 6. Perfect shuffle based Benes network modified to include pre- and post-shuffle stages.

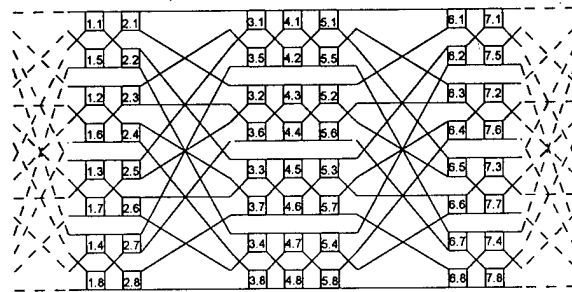


Figure 7. Higher order shuffle Benes network topologically equivalent to Figure 6.

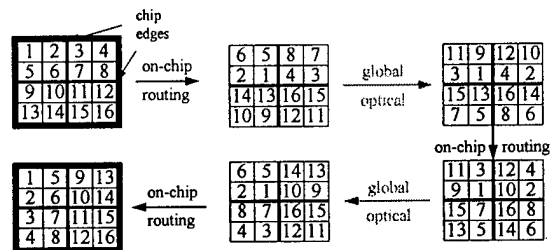


Figure 8. Example Two-Bounce interconnection pattern: Transpose.

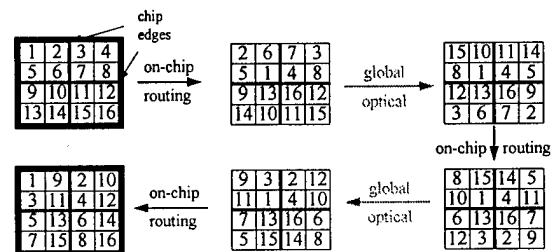


Figure 9. Example Two-Bounce interconnection pattern: Folded Perfect Shuffle.

optical system into the single plane Two-Bounce module. Additionally, the Two-Bounce architecture requires one lens per chip, so a Two-Bounce module performing a Folded Perfect Shuffle on a 4×4 lens array, would utilize 16 lenses and perform unity magnification.

The Two-Bounce generalizes directly to any permutation network of size $N=2^M$, where M is even. For example, if $M=10$, $N=1024$, then an optimum choice for k is $2^{M/2} = 32$. Therefore at least 32 lenses are required for the reflective optical shuffle module. These lenses could be arranged in a 4×8 pattern, or more lenses could be utilized to make the array square. For networks of arbitrary sizes, two approaches can be considered. The network can be mapped onto the next largest readily packaged size array (2^M where M is even) or if the interconnection can be partitioned into a number of separate smaller arbitrary permutations, then each of these can be interleaved and implemented in parallel with a single optical system.

It has been pointed out that a symmetric k -shuffle network ($k=N^{1/2}$) allows any node to communicate with any other node with a single pass through the optical system and 2 stages of switching [21]. This is a k -shuffle based banyan, and therefore suffers from internal blocking (For a full permutation this amounts to $\sim 2/3$ of the data [22].), i.e., not every node can simultaneously communicate with another node – not all permutations are possible. Since the Two-Bounce architecture implements the full Benes network, it truly achieves arbitrary rearrangeably non-blocking performance with just *two* optical passes. These two optical passes provide the necessary global interconnection for the entire Benes network – all other interconnections are local and therefore are contained within each chip.

III. Optomechanical Evaluation System

In order to demonstrate the functionality of the Two-Bounce architecture, A 16 node, 4 "chip" module was designed and built. The purpose of this module was to effect the Viterbi Decoding trellis (simultaneous forward and backward perfect shuffles), but it also validates the Two-Bounce concept [15]. Since OEICs with the required functionality are not currently available, a fiber-coupled array was utilized for optical input and output. This Two-Bounce prototype has 64 fibers mounted in a faceplate. The capacity of the Two-Bounce module greatly exceeded the number of fibers utilized for the architecture verification. Figure 10 shows the system demonstration prototype experiment. A laptop PC performs the requisite smart pixel functionality and is interfaced through a data acquisition system to an emitter/detector driver box. This box provides the electronic interface for the fiber-coupled arrays in the optical prototype. If smart pixels were utilized only the small optical module would be present in this system. The rest of the hardware is used to mimic the smart pixel functionality. The optical interconnection module is shown in Figure 11. The central component of the optical interconnection module is the 2×2 lens system, designed to interconnect a 2×2 array of smart pixel ICs. We are investigating custom designed lenses for the prototype, but the initial system uses commercially available miniature projection lenses that have a wide and flat field-of-view, with high resolution. Using a VCSEL array as an input source, we characterized several low f-number lenses that are commercially available. The selected lenses are f/1.1, with useable fields of approximately .8 cm across, and with resolution spot sizes of approximately $10 \mu\text{m}$ – consistent with the anticipated parameters of smart pixel integrated circuits. The 4 lenses in the array were selected from a larger set to match their parameters to a high degree. An active alignment procedure was developed for the module that involved individual positioning of the lenses over the smart pixel backplane. This procedure has demonstrated a registration accuracy of ~ 10 micrometers over the backplane, as large as 10 cm and lens arrays as large as 4×4 , utilized for the optical prototype [10], which is consistent with the anticipated required smart pixel IC alignment dictated by sources such as VCSELS that will be 10-20 micrometers in diameter.

Figure 11 shows the optomechanical prototype module. It consists of 3 planes on a group of 4 rails. One plane holds the thin fiber I/O plane, simulating the smart pixel optical I/O. The fiber

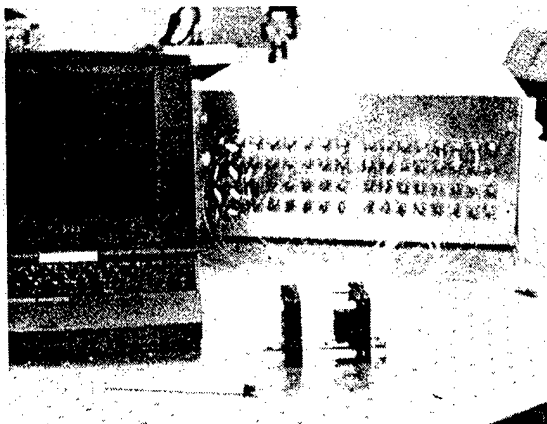


Figure 10. Two-Bounce Demonstration System

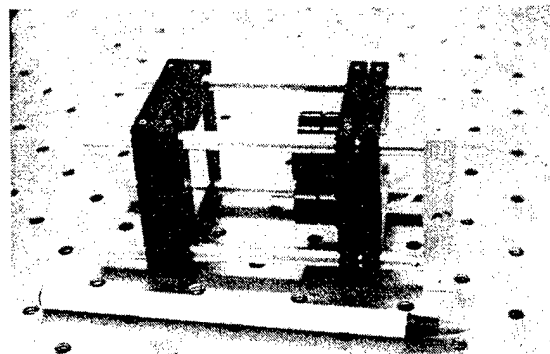


Figure 11. Close up view of optical interconnection module.

plane was precision machined to mount the 64 fibers at the desired locations across the smart pixel plane. The second plane is the lens support plane; it positions each lens over its corresponding optical I/O and maintains the alignment between the two. The third plane holds the mirror for the retro-reflective interconnection. These three planes must be able to be adjusted so that they are perpendicular to one another within a few degrees. Since this parallelism is set *before* the lens alignment is performed, it only effects the efficiency of the interconnect, not its alignment. If the mirror was not quite perpendicular to the smart pixel plane, the lenses would be aligned with this error present and the proper interconnects would still be achieved. The parallel and symmetrical nature of the optical interconnect provides some cancellation of distortion effects, as it is a reciprocal optical system. The system has been designed to receive actual packaged smart pixel devices (in place of the fiber-coupled plate), as they become available. As emitter-based smart pixel technology rapidly matures, i.e., as the pixels get "smarter" (more integrated digital logic) and have higher densities of optoelectronic I/O, the Two-Bounce prototype will be able to readily incorporate them. The optical module was comprised of groups of I/O sites totaling 32 emitters and 32 receivers interconnected in a shuffle pattern. The overall dimensions of the system are approximately 4 cm \times 4 cm \times 8 cm.

A photograph of the experimental set up used to evaluate the Two-Bounce module is shown in Figure 12. The experimental system is comprised of three planes: the OE I/O plane, the lens

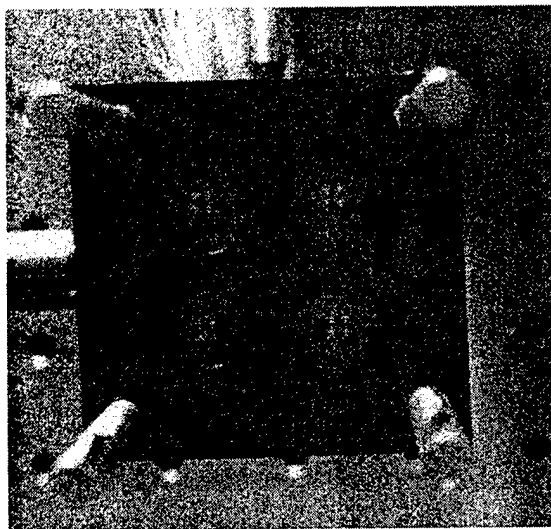


Figure 12. Photograph of fiber coupled simulated smart pixel I/O plane in experimental setup (mirror removed).

plane and a mirror plane. The OE I/O plane holds the fiber array in place of a smart pixel OEIC. The experiments discussed in the next section replaced this plane with VCSEL arrays and a CCD imaging system. These planes are supported so that the inter-plane distance and parallelism may be adjusted. In the experiments, the backplane was used as the reference plane to which all other elements in the system were aligned. The lens array plane is comprised of a flat aluminum plate with 4 apertures, one for each lens in the lens array. A lens was precisely aligned above each group of I/O sites (OEIC) using a self-alignment procedure that is amenable to automation [10].

IV. Scaling Laws for Free-space Optical Interconnection Systems

A. Motivation

“Smart pixel” throughput capabilities are projected to exceed 1 Tbit/s/cm² [23]. The hope is that this capacity will enable free-space optical interconnects (FSOI) to provide significant throughput, size, and power consumption advantages over all-electronic interconnection technologies. To accomplish this goal, new architectures for interconnection-limited problems must be devised which exploit the ability of smart pixels to combine parallel high density I/O with local electronic logic.

Clearly, optics provides the highest potential payoff for those problems that must dedicate a large amount of resources to interconnecting multiple processors in a dense, compact environment that challenges conventional electrical interconnect packaging approaches. In particular, 3-D free-space optics may offer the ability to overcome the throughput and global interconnection limitations of conventional 2-D metallic interconnection technology by exploiting the additional spatial dimension. The purpose of this paper is to explore and compare the geometric scaling rules for 2-D metallic and 3-D free-space optical interconnection topologies. Such scaling relationships will be useful in quantifying the benefits of optical interconnection approaches in given problem domains.

The focus of this paper is on those multi-processor applications that require global high density interconnections characterized by high minimum bisection bandwidth (BB) – a widely accepted measure of the degree of interconnection difficulty in networks. The BB of a network is defined as the bandwidth that crosses a boundary that cuts the network in half – it is a measure of wiring difficulty [17]. In architecture design, there is a direct trade-off between minimum BB and latency in a network. It is therefore generally desirable to implement networks with the largest minimum BB that can be practically achieved to solve a given problem. The ability of optical elements to interconnect large arrays in space-variant patterns, without crosstalk in the medium, suggests that FSOI techniques are particularly promising for problems with high BB. In particular, optical space-variant approaches to performing high BB perfect shuffle [3] and related patterns have been studied for some time. [24, 4-8]

Chip area requirements for high density, interconnection – limited integrated circuits were found to be proportional to BB^2 . [25]. In this paper, circuit area analysis is extended to problems for which the integrated circuit (IC) interconnection area is not sufficient to achieve the desired multiprocessor links. The total interconnection area must therefore be determined for interconnection packaging technologies lower in the interconnection hierarchy, i.e., for multi-chip modules (MCMs) and – for the most highly interconnected problems – for printed circuit boards (PCBs). The total area requirement is used as a basis for estimating performance costs, such as volume, latency, and power consumption.

In general, the total circuit area will be the sum of interconnection area and processor area. The focus of this paper is on those problems for which the area dedicated to inter-processor interconnection dominates. It follows that the volume, latency, and power consumption performance metrics will then also be limited by the interconnection requirements.

In Section 2 of this paper, basic VLSI electrical interconnection area scaling requirements are extended to MCMs and PCBs, and then extended further, to latency, power, and volume scaling rules. These parameters are derived as a function of BB of the architecture. Section 3 is a

derivation of the same parameters for interconnections based on opto-electronic technology. The emphasis is on globally interconnected systems, in which multi-chip data interchange dominates the interconnection requirements. In the discussion of Section 4 the derived scaling laws for different interconnection technologies (electrical, macro-optical, micro-optical) are compared to define those problem domains in which each technology has the greatest benefit. The Conclusion, Section 5, summarizes the key results and relates the analysis to recent experimental developments.

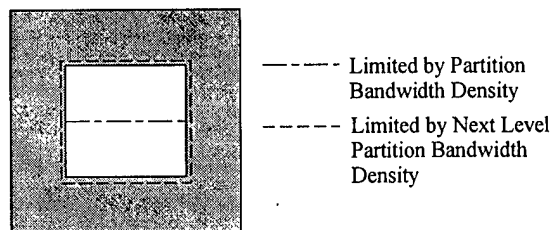


Figure 13. Internal and external interconnection bandwidth partition boundaries for metallic interconnection technologies.

B. Electrical Interconnection Requirements

Network Partitioning

The starting point for performance scaling analysis is the bandwidth (BW) density capacity of the interconnection technologies. For the electronic packaging hierarchy the linear BW density is different in each level of the packaging hierarchy (IC, MCM, PCB). The linear BW [measured in Terabits/s/cm] density stipulates the maximum bandwidth that can cross any boundary as a function of the length of the boundary. Two types of boundaries readily lend themselves to this analysis: internal bisection boundaries within partitions, and external boundaries between partitions. Figure 13 depicts these two types of bandwidth-limited boundaries for the case of a single IC partition placed on an MCM.

In order to relate linear BW density to area requirements, the architecture is repeatedly partitioned into smaller equally sized sets of nodes. The requirements of every sub-partition are calculated based on the linear BW density of the interconnection technology. Often the optimum partitioning of the system – in the least area sense – is the minimum bisection that separates the network into two equal groups and “cuts” the fewest “wires”. However, in general partitioning into any prime number of groups should be considered. For example, it is possible that the optimum partition – one that minimizes the bandwidth between partitioned subsets – of a group of nodes might be a tri-section (three equal sized groups of nodes with less “wires” cut than a bisection of the nodes). To simplify the discussion, bisection partitions are assumed in this paper.

Figure 14 depicts an example interconnection architecture with 16 nodes. The I/O requirement for the entire system is 8 B, where B is the bandwidth of a single “wire” and there are four inputs and four outputs. Figure 15 depicts the minimum BB partitioning of the system, the “cut wires” are depicted as dashed lines. The internal minimum BB of the architecture is seen to be 6 B. These “cut wires” are now part of the external bandwidth requirements and are therefore added to I/O bandwidth requirements. Figure 16 depicts the next level of minimum bisection partitioning. In general, this partitioning is repeated until each partition contains only one node. Figure 17 is a tree depicting the resulting partitions for the example network shown in Figures 14-16. Each node of the tree is labeled with the partition, the bisection bandwidth requirements of the partition, and the I/O requirements of the partition. Network partition trees

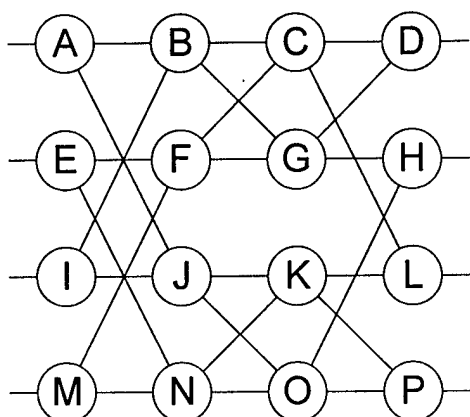


Figure 14. An example interconnection network of 16 nodes.

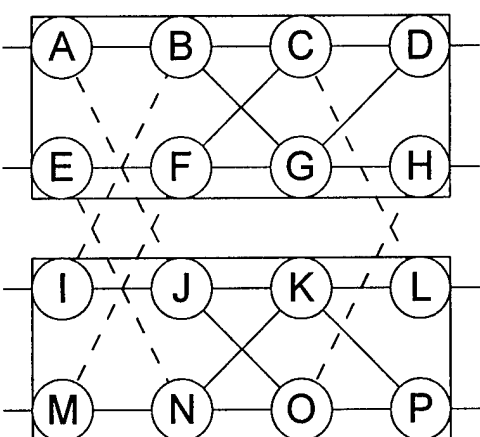


Figure 15. Top level minimum bisection partitioning of the nodes depicted in Figure 14.

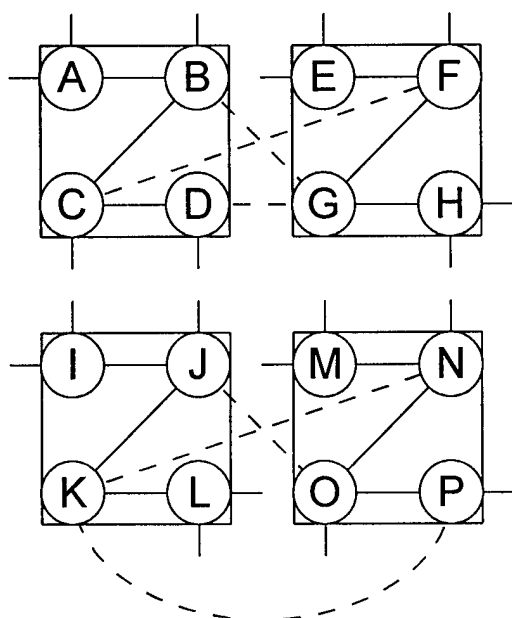


Figure 16. Sub-partitioning of Figure 3 into second level bisections.

are useful in determining the requirements for the interconnected architecture in different technologies.

To relate the BB and I/O requirements of a node of the tree to area, the maximum capacities of the different levels of the packaging hierarchy must be determined. This is driven by the maximum practical or realizable size of each level. If one assumes a maximum size of a square package ($A^{1/2}$), with uniformly distributed nodes, and a linear bandwidth density (D_{layer}) for that layer, then Equations 3 and 4 dictate the maximum partition BB and I/O of that packaging layer:

$$BB_{\text{max}} = \sqrt{A_{\text{max}}} D_{\text{layer}}, \quad (3)$$

$$IO_{\text{max}} = 4\sqrt{A_{\text{max}}} D_{\text{layer}}. \quad (4)$$

It should be noted that when the partition boundary coincides with a technology boundary, e.g., the partition is an entire chip placed on an MCM, the I/O D_{layer} is determined by the lower hierarchical layer. As illustrated in Figure 13, all data lines that leave a chip must cross the chip package perimeter in the MCM layer, no matter how dense the connections between the chip and MCM.

When the maximum capacities of each layer for bisection bandwidth and I/O bandwidth are known, the tree of Figure 17 can be traversed to calculate the required substrate area. Beginning with any node in the bottom row, determine first if that node can be realized within a single IC, and if so, then determine what size is required. If it can be realized, then traverse up the tree to its parent node. Now it must be determined if the parent node is realizable, while simultaneously realizing both of its daughter nodes in half an IC. The tree is climbed in this fashion until a given partition cannot be realized in the IC layer. From this point the process continues using lower packaging layers (e.g., MCM followed by PCB) until the root node is reached. When this node is reached the total interconnection substrate area is estimated by calculating the maximum total area required

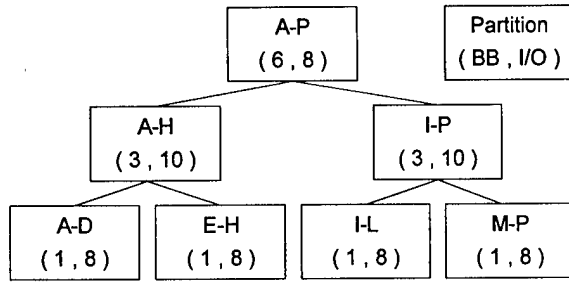


Figure 17. Bisection tree of depicting two levels of bisection of Figure 2. Each node is labeled with the partition name and the internal BB of the partition and its I/O bandwidth requirement. This tree can be

across all three layers of the hierarchy. Note that the area specified by the bisection tree is only the area required for interconnection. It is possible that the total chip area, and not the lowest hierarchical layer, i.e., the topmost bisection, drive the area requirements. For example, when this analysis is applied to an architecture in which the first BB (topmost node) is extremely low, but the subsequent partitions are characterized by large BB, the interconnection area requirement of the topmost node (e.g., PCB) will not result in a area large enough to mount the resultant MCMs and ICs. In this case, the higher

packaging layers clearly drive the interconnection area requirements. This system would be characterized by having many dense ICs interconnected on MCMs, but with little interconnection between the MCMs in the PCB layer. In this case, the maximum of the MCM or IC area would determine the overall architecture area requirement.

When a network requires a “regular global interconnection pattern,” defined as architecture for which each level of the bisection tree results in half the BB requirements of the previous level, the topmost partition determines the overall area. Butterflies and shuffles are examples of regular global interconnection patterns. In this case, the above analysis is a direct extension the VLSI area complexity analysis [25] to lower levels of the hierarchy. When the architecture is not a regular global interconnection pattern, the first bisection does not necessarily drive the area requirements. In this case, the bisection tree provides a mechanism to identify and quantify the area driving interconnection bottlenecks of the architecture. The following section extends this analysis to equations for area, power, latency, and volume.

Geometric Scaling Rules for Planar Metallic Interconnections.

The previous section identified the optimum partitioning of a network and determined the BB requirements for each partition. From this, the area required for each partition is given by:

$$A_{i,j} = \text{MAX} \left[\left(\frac{BB}{D_{\text{layer}}} \right)^2, A_{i-1,2j} + A_{i-1,2j-1} \right], \quad (5)$$

where i is the layer of the tree (numbered from bottom to top in Figure 17) and j is the node within that layer (numbered from left to right). This equation states that the interconnection area requirement of a node is the maximum of its own BB requirements and the sum of its two daughter nodes' requirements.

The substrate area interconnection requirement can be used to determine other important performance parameters. For example inter-processor signal latency may an issue when synchronous operation of the multiple processors is desired. In planar metallic technology the worst case maximum path length, L_{max} , between processors will be the diagonal distance across the interconnection substrate:

$$L_{\text{max}} = \sqrt{2A} = \sqrt{2} \frac{BB}{D_{\text{layer}}}, \quad (6)$$

where A is the area requirement. To the extent that latency is proportional to the maximum distance between processors, L_{\max} is a measure of latency in the network.

The total packaging volume for the interconnection can be bounded by assuming that each layer of the metallic interconnection hierarchy has a finite height, H_{layer} , that is the required clearance for the enclosure of the circuit, as determined by practical packaging constraints. For example, possible enclosure heights for the three levels of metallic packaging might be 0.1, 0.5, and 1 centimeters for H_{IC} , H_{MCM} , and H_{PCB} , respectively. The volume required for a given metallic interconnection package is therefore:

$$V_{\text{layer}} = H_{\text{layer}} A = H_{\text{layer}} \left(\frac{BB}{D_{\text{layer}}} \right)^2. \quad (7)$$

The interconnection network's power consumption requirement is also related to the geometric constraints of the planar interconnection hierarchy. Although the exact scaling rules for power consumption will depend on the details of the metallic technology used and other operational characteristics, it is useful to bound the power requirement scaling rules for later comparison with optical interconnection scaling rules. If the electrical interconnections within a level are viewed as lumped capacitive loads (as for example in the short interconnections on ICs), then the power will scale as the average length of a line. In this domain the power requirements are bounded by:

$$P_{\text{capacitive}} = P_c \sqrt{A} D_{\text{layer}} \sqrt{A} = P_c D_{\text{layer}} A = P_c \frac{BB^2}{D_{\text{layer}}}, \quad (8)$$

where P_c is the power required per unit length per unit bandwidth [W/cm/THz], so the product of P_c and D_{layer} has units of Watts/cm². This represents an upper bound on the power requirements. A lower bound is derived under the assumptions of lossless transmission lines for the propagation of data. In this case, the power is bounded from below by:

$$P_{\text{lossless}} = P_l BB, \quad (9)$$

where P_l is the power required [W/THz] to drive the lossless transmission lines. Equations 8 and 9 provide bounds on the trends of power scaling as a function of BB in the metallic packaging hierarchy. An actual implementation will therefore likely scale somewhere between the lower bound, which scales as BB , and the upper bound, which scales as BB^2 . These bounds are presented here to facilitate a later comparison with optical interconnection requirements.

C. Optical Interconnection Requirements

Representations of Free Space Optical Interconnections

FSOI based systems can be categorized by the ratio of lenses to optical I/O. This is a measure of the degree of space variance in the optical system. Figure 18 is a depiction of the range of FSOI approaches. In general, planes of optical I/O may be interconnected to each other. For simplicity, Figure 18 depicts single plane reflective architectures in which all of the smart pixel resources are distributed on a common plane. Figure 18a depicts a one chip per one lens scheme, termed a *macro*-optical interconnection because the lenses are approximately the size of the smart pixel chips – several millimeters or larger [10]. In this case, many optical I/O are located beneath each lens. Figure 18c depicts a *micro*-optical approach with one lens for each optical

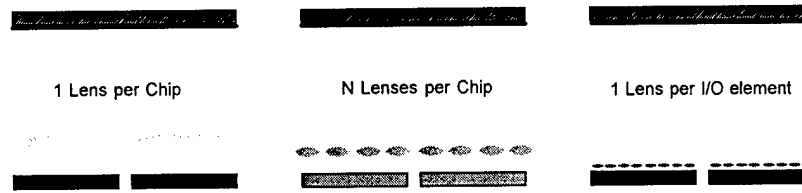


Figure 18. The three basic optical interconnection approaches, shown for a reflective architecture.

I/O. In this case the beam steering elements have diameters equal to the pitch of the high density opto-electronic I/O – on the order of 100s of microns. Figure 18b depicts an intermediary approach with several lenses per chip, and several I/O per lens. In principal, a single lens per I/O provides the maximum flexibility, since an arbitrary interconnection pattern can be implemented with appropriate lens elements. To achieve arbitrary interconnection patterns in an approach like that of Figures 18a and 18b requires local electronic interconnections and multiple passes through the optical system [12]. The shape of the modules depicted in Figure 18 is dependant upon the $f_\#$ of the optics utilized. As the $f_\#$ approaches unity, the reflective module approximates a cube in form [1,13].

As depicted in Figure 18, the interconnection architecture consists of an array of point-to-point links. In principle, scaling to larger arrays, with larger BB, simply requires larger multi-chip smart pixel arrays with the interconnection volume scaled appropriately to maintain the approximately cubic aspect ratio. Such scaling will entail longer link lengths. Under the assumption that smart pixel based interconnections will require opto-electronic densities of $\sim 1000/\text{cm}^2$, then diffraction losses will limit the lengths of these links. Macro-optics, with lens sizes of mm's or more, scale well into free space volumes with sizes of 1000s of cm^3 . However, the combination of high I/O density and long link paths will lead to diffraction limits in the micro optical approach and thereby affects the scaling properties.

To determine the performance scaling of micro optics requires a determination of the maximum allowable throw distance between optical elements above emitters and detectors. The optical element size is set by the pitch of the smart pixel I/O, limiting it to 100's of μm 's or less. Assuming a Vertical Cavity Surface Emitting Lasers (VCSELs) for the optical emitters, and 200 μm optical elements, the propagation of Gaussian beams can be applied [26-29]. The loss and crosstalk tolerances of the design and the type of beamforming that is implemented determine the actual throw distance. For example, the micro-elements can be configured to achieve minimum divergence or minimum beam waist. Both yield similar throw distance results. The following example is a minimum divergence angle estimate. The loss criteria are set as follows: the input lens should capture at least 99.9% of the VCSEL light (to allow a close approximation to Gaussian beam propagation between the two lenses), and the throw distance should be constrained by the requirement that the receiving lens capture 86% of the light (i.e., matched to the beam waist). Given a micro optical element with diameter d , focal length f , and VCSEL beam waist w_0 , the beam waist at the transmitting element aperture is given by:

$$w_1 = w_0 \sqrt{1 + \left(\frac{f\lambda}{\pi w_0^2} \right)^2} = \frac{d}{2k}, \quad (10)$$

Where λ is the wavelength and k is chosen as 2.12 to maintain the Gaussian approximation to collect 99.9% of the light at the transmitting aperture [27]. The beam waist at the receiving element is therefore given by:

$$w_2 = w_1 \sqrt{1 + \left(\frac{z_{\max} \lambda}{\pi w_1^2} \right)^2} = \frac{d}{2}, \quad (11)$$

Equations 10 and 11 can be solved to determine the maximum throw distance z_{\max} :

$$z_{\max} = \sqrt{k^2 - 1} \left[\frac{\pi d^2}{4 \lambda k^2} \right]. \quad (12)$$

As an example, with $d=200 \mu\text{m}$, $\lambda = 0.85 \mu\text{m}$, and $k=2.12$, the maximum throw distance equals $\sim 1.5 \text{ cm}$. This first order approximation assumes that the beams propagate perpendicular to the optical elements. The throw distance will actually be reduced for optical beams that propagate at steep angles due to the cosine projection of the beam waist.

From geometric constraints, z_{\max} and the $f_{\#}$ of the optics determine the mirror height (h), as given by:

$$h = \frac{f_{\#} z_{\max}}{\sqrt{1 + 4 f_{\#}^2}} \quad (13)$$

Diffraction effects on micro optical architectures dictate that large BB systems, characterized by large substrate areas, do not retain the cubic form of macro-optics. The short distances of micro optics dictate a low aspect ratio for the interconnection volume. Furthermore, this short throw distance limits the lateral displacement of any given link, thereby requiring repeaters to connect globally distributed nodes. This need for repeaters greatly impacts the scaling of micro optical architectures as detailed below.

Geometric Scaling Rules for 3-D Smart Pixel Based Architectures

Since FSOI interconnections are not confined to planar links, the interconnection density limitations stem from the *area* I/O density capabilities of smart pixel technology and the ability of optical elements to perform the inter-chip data interchange functions. FSOI concepts based on interleaved imaging of sub-arrays, such as depicted in Figure 18a and 18b, are able to link arrays of smart pixel I/O with resolution well beyond that required to achieve the anticipated Tbit/sec/cm² I/O densities of smart pixel arrays. The area of the smart pixel surface and the density $D_{I/O}$ (Terabit/s/cm²) of the optical I/O therefore determines the maximum bandwidth crossing external boundaries for FSOI. If the interconnection pattern is global, in that every IC communicates with every other IC, then half of the total IC area contains optical I/O which cross any bisection boundary. The BB capability of FSOI is thus given by 1/2 the total smart pixel I/O. For example, if $D_{I/O} = 1 \text{ Tbit/sec/cm}^2$, then the bisection bandwidth capability of FSOI is $D_{I/O} A_c / 2$, where A_c is the total smart pixel chip area employed. Inverting this, the area required for macro optical interconnections is:

$$A_{\text{Macro}} = \frac{2BB}{D_{I/O}}. \quad (14)$$

As discussed previously, micro optic's requirement for repeaters changes this for multi chip architectures by reducing the effective density of I/O, including only those emitters originating and not repeating data. The equation for micro optics in terms of this effective density (D_{eff}) is:

$$A_{micro} = \frac{2BB}{D_{eff}}. \quad (15)$$

Where D_{eff} is given by:

$$D_{eff} = \frac{D_{I/O}h}{f_{\#}\sqrt{A_{micro}}}, \quad (16)$$

where $h/f_{\#}$ is a normalized lateral throw distance, and A_{micro} is the new area requirement. Solving for A_{micro} in equations 15 and 16 yields:

$$A_{micro} = \frac{4BB^2 f_{\#}^2}{D_{I/O}^2 h^2}. \quad (17)$$

Volume, latency, and power requirements may be derived directly from the above area analysis. Since there is no packaging hierarchy in free space optics, only one area bandwidth density is required. The volume required for macro optical systems is approximated by:

$$V_{Macro} = A_{Macro}^{3/2} f_{\#}, \quad (18)$$

whereas the fixed throw distance of micro optics results in a volume of:

$$V_{micro} = hA_{micro}. \quad (19)$$

From geometry, the maximum path length for both macro- and micro optical architectures is:

$$L_{max} = \sqrt{A(1 + 2f_{\#}^2)}. \quad (20)$$

However, the area for micro- and macro- optical systems scales differently resulting in a different overall scaling in maximum path length. Similarly, the power requirements of optically interconnected modules derive directly from area requirements. These requirements are given by:

$$P = ANP_{link}. \quad (21)$$

where A is the area populated by I/O, N is the total number of I/O per cm^2 , and P_{link} is the power per I/O link.

D. Discussion

Tables 1 & 2 contain example parameters to make comparisons between planar metallic, micro optical, and macro optical interconnections. While the actual values may vary, the slopes of the scaling equations are fixed. Figure 19 plots the area scaling of planar interconnects, micro optical interconnects, and macro optical interconnects based on the sample parameters. Figure 19 shows the FSOI area requirement grows in direct proportion to the BB requirement [30]. However, this scaling argument applies only to the macro optical architecture. The macro and micro optical architecture scale identically until the micro optical architecture hits its diffraction limited throw distance (at < 1 Tbits/sec). At higher BB, the micro optical architectures scale at the same rate as the metallic architectures

Table 1. Example Planar Interconnect Parameters

	IC	MCM	PCB
A_{\max}	6.25 cm^2	225 cm^2	N/A
D_{layer}	0.5 Tbit/cm	0.2 Tb/cm	0.1 Tb/cm
H_{layer}	0.1 cm	0.5 cm	1 cm
P_c	5 W/cm^2	5 W/cm^2	5 W/cm^2
D_{layer}			
P_{lossless}	N/A	20 mW/Gbit	40 mW/Gbit

Table 2. Example Optical Interconnect Parameters.

	Macro Optical Interconnects	Micro Optical Interconnects
Wavelength	850 nm	850 nm
H_{\max}	N/A	1 cm
$f_{\#}$	1	1
$D_{\text{I/O}}$	1 Tbit/s/cm^2	1 Tbit/s/cm^2
P_{link}	5 mW/link	5 mW/link
N	1000 links/cm^2	1000 links/cm^2

Figure 20 depicts the interconnection volume scaling requirements for the discussed technologies. Note that while micro optics has a much larger area than macro optics, the difference in volume is not as extreme. This is due to the form of the micro optical and macro optical architecture. Micro optical architectures are broad and flat, whereas macro optical architectures are cubic in nature. However, after the diffraction limited throw distance is exceeded the micro optical volume requirements scale as BB^2 , as does electronics, whereas the macro optical volume scales only as $BB^{3/2}$. For the selected parameters, both micro optical and macro optical architectures have 2 or more orders of magnitude over metallic approaches for BB approaching 10 Tbits/sec. It noteworthy that the “apparent” wasted volume, that cubic shaped optical interconnection architecture seems to have, actually leads to this significant advantage.

Figure 21 depicts the maximum path length scaling requirements for the discussed technologies. Clearly for “low” BB ($< 1 \text{ Tbits/sec}$), IC technology is superior. However, for greater bisection bandwidths, macro optical path lengths scale as $BB^{1/2}$, whereas micro optical and electronic path lengths all scale linearly with BB. As discussed before, the maximum path length relates directly to the latency and skew in the synchronization of multi-processor systems. The data show that macro optical systems will have a significant advantage in latency in the $\sim 10 \text{ Tbits/sec}$ BB regime.

Finally, Figure 22 depicts trends for the interconnection power consumption requirements for the relevant technologies. The electronic packaging layers are bounded on the graph by lossless transmission line analysis and lumped capacitive loading. Note there are two lines each for MCM and PCB layers representing these bounds. Macro optics again achieves the best scaling ($\sim BB$) and matches the best possible electronic scaling slope. Micro optics, however, scale as poorly as the worst case electronic power requirements ($\sim BB^2$).

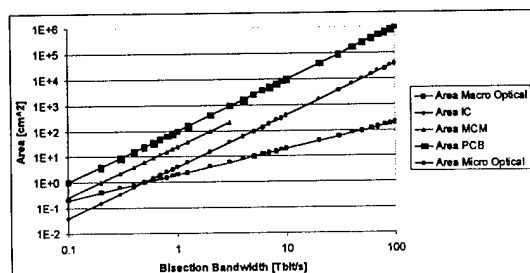


Figure 19. Area scaling graph for macro optical, micro optical, and metallic planar interconnections.

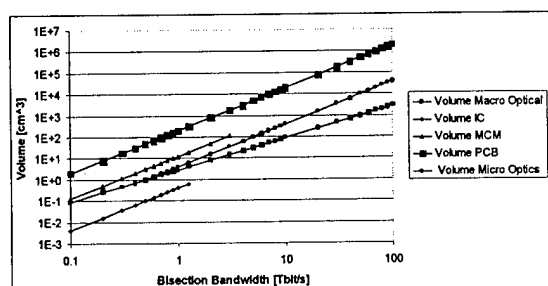


Figure 20. Volume scaling graph for macro optical, micro optical, and metallic planar interconnections.

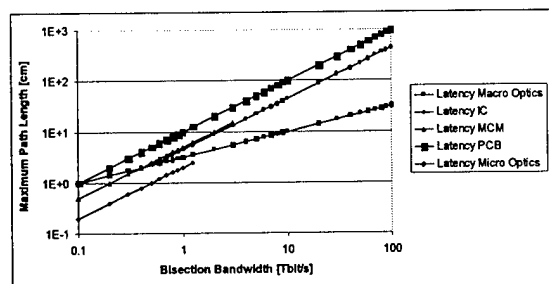


Figure 21. Maximum path length scaling graph for macro optical, micro optical, and metallic planar interconnections.

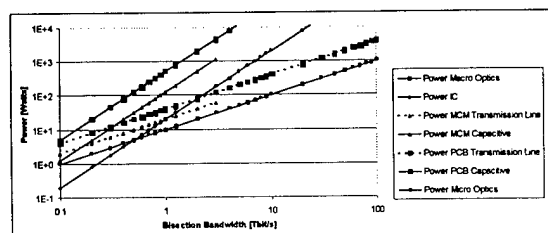


Figure 22. Power scaling graph for macro optical, micro optical, and metallic planar interconnections.

Although all of the performance metrics described above are derived from substrate area considerations, user defined metrics may combine them. For example the product of power consumption and volume may be a critical figure of merit for some applications. As can be seen from Figures 21 and 22, the advantages of macro-optical FSOI architectures are amplified when such measures are combined.

To realize the potential of the rapid advances being made in high throughput smart pixel technology, architectures based on macro-optical interconnection modules must be developed. Figure 23 is a photograph of a prototype macro-optical reflective multi-chip interconnection module. This system links 4 smart pixel chips with a 2x2 array of miniature projection lenses. This approach has achieved accuracies of $\sim 10 \mu\text{m}$ across MCM substrates of 10 cm in extent for lens arrays as large as 4x4 [10].

The fundamental conclusion of this analysis is that FSOI approaches have the most favorable scaling advantages when multiple ICs are globally interconnected – i.e., when multiple chips are communicating simultaneously with multiple chips. This scenario typifies the multi-Terabit/sec BB regime in which FSOI has the greatest payoff. The fundamental advantage of macro-optical FSOI over metallic interconnections, in terms of substrate area based metrics, does not rely on the actual bandwidth densities of the routing layers. It stems directly from the *reduction* in density in metallic interconnections as bandwidth is placed in lower layers. The only technological improvement that would overcome this fundamental advantage is if the lowest routing level (PCB) densities approached the densities of optical interconnections. This is not projected to happen, as density increases tend to "trickle down" from increased chip densities to increased MCM densities, to increased PCB densities. As long as the metallic packaging

hierarchy remains the advantage of FSOI will hold true. In other words – although electronic interconnection technology will continue to improve in density (as, we hope, will smart pixel-based FSOI technology), the height and placement of jumps between the metallic interconnect curves will change somewhat. However the basic and fundamental advantage of FSOI, as embodied in the lower slope and lack of partition boundaries (i.e., *no interconnect packaging hierarchy*) for the optics will remain.

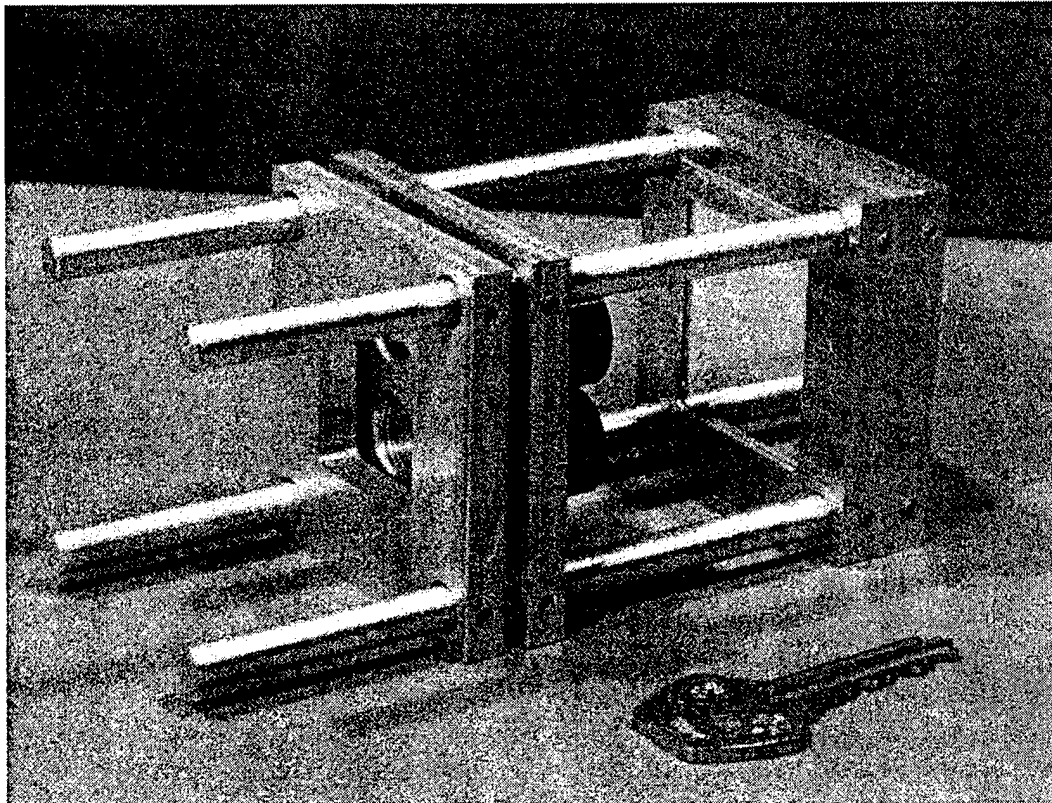


Figure 23. A photograph of a prototype macro-optical reflective multi-chip interconnection module.

V. Hybrid Macro/Micro-optical Interconnection Concept

Free-space optical interconnections are projected to provide bandwidth densities on the order of a Terabit/sec/cm² [31]. Scaleable multi-terabit interconnection fabrics may be achieved using multiple optoelectronic integrated circuits linked to each other in a global high bisection bandwidth pattern [2] as depicted in Figure 1. In this configuration each lens links the optical I/O from a single chip, located at the lens' focal plane, to all chips in the receiving array. Clusters of emitters, such as vertical cavity surface emitting lasers (VCSELs), and detectors are imaged onto corresponding clusters on other chips such that many point-to-point links are established in an interleaved optical shuffle pattern across the multi-chip plane. Monolithically integrated VCSEL/detector arrays, with emitter and receiver elements of 10 and 50 μm , respectively, and with element-to-element spacing as small as 100 micrometers, have been evaluated in a prototype shuffle system [32]. With such I/O density and pitch, the global optical interconnection module must provide flat, high resolution, near distortion-free image fields, across a wide range of ray angles in order to avoid cross-talk and maintain high link efficiency.

Although modern optical design and manufacture techniques provide approaches to achieving high resolution, registration accuracy is more problematic. Registration accuracy may be defined as the difference between the location of the image of a VCSEL and the location of its corresponding detector. Registration must be maintained at a level less than the size of the detector ($\sim 50 \mu\text{m}$) across the entire multi-chip plane ($\sim 10 \text{ cm}$ wide). Distortion in the optical system will cause poor registration performance in the system. It is well known that holosymmetric systems (systems with radial symmetry about their optical axis and symmetry along their optical axis about their aperture) cancel distortion [33-35]. While the interconnection system depicted in Figure 24 appears to be symmetric, the aperture of the system is not at the midpoint between the transmitting and receiving lens planes. As depicted in Figure 25a, this asymmetry results from the normal orientation of the VCSEL beams – parallel to the optical axis. In order to cancel distortion, the effective aperture must be moved to the midpoint between the transmitting lens and receiving lens. Unfortunately, placing the aperture at this location causes the narrow VCSEL beams to miss the aperture entirely or be severely vignetted. This vignetting can be corrected, if the VCSELs are steered to emit at angles that cause them to propagate through the new central aperture as shown in Figure 25b. This is possible only because the VCSELs have narrow beam divergence. Once the VCSELs have been steered through the central aperture no physical aperture is needed at this location. The proposed method for implementing the beam steering is depicted in 25c. A linear diffraction grating or prism is placed above each VCSEL and detector. In this configuration, each VCSEL's beam is deflected by an angle which causes its beam to cross the optical axis at the halfway-point between the transmitting and receiving lenses. To maintain symmetry, and hence eliminate distortion, identical micro-elements must be employed at the detector plane as well, as depicted in Figure 25c.

Figure 26 shows the deflection angle, ϕ , as it relates to the geometry of the other variables of the interconnect system for the on-axis cluster. The off-axis distance of the VCSEL under consideration is x , the focal length of the lens is f , $f\#$ is the ratio of this focal length to the lens diameter, θ is the angle of the collimated beam with respect to the optical axis from the VCSEL, N is the number of chips on one side of the square array (see Fig. 24), x_L is the height the

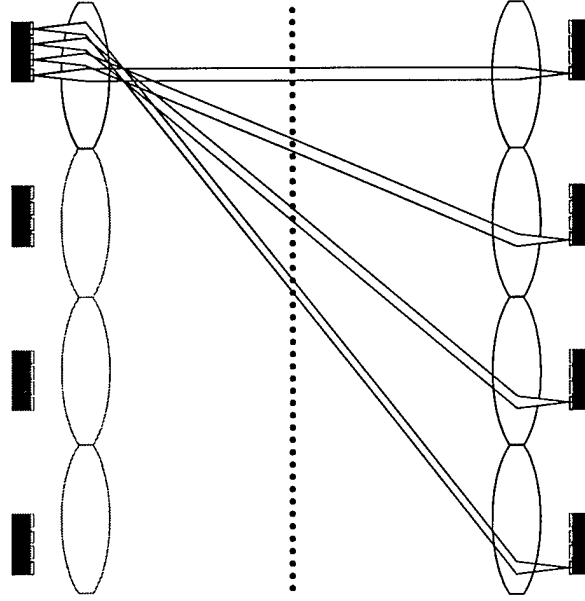


Figure 24. Schematic side view of global optical shuffle interconnection. There is one lens over each chip. Each chip communicates with every chip in the receiving array. The system can be folded along the dotted line to facilitate packaging and alignment.

deflected beam hits the lens plane, d is the distance from the VCSEL plane to the diffraction grating, and Δx is the effective displacement of a VCSEL emitting parallel to the optical axis.

In Figure 26 there are two congruent relationships:

$$\tan \theta = \frac{x}{f} = \frac{2x_L f_{\#}}{Nf}, \text{ and} \quad (22)$$

$$\tan \phi = \frac{\Delta x}{d} = \frac{(x_L - x)}{f}. \quad (23)$$

From Equations 22 and 23, the deflection angle as a function of x the following is:

$$\phi = \arctan \left[\frac{x}{f} \left(\frac{N}{2f_{\#}} - 1 \right) \right]. \quad (24)$$

Figure 27 demonstrates that as x varies along the cluster the deflection angle varies in such a way as to make the collection of prisms or gratings act as a negative lens. The focal length (f_{eff}) of this effective lens is given by:

$$f_{\text{eff}} = \frac{f}{\frac{N}{2f_{\#}} - 1}. \quad (25)$$

The above analysis can be extended to the general multi-chip and off-axis case for inter-chip connections in Fig. 24. In this case the aperture remains at the midpoint between the two lenses, but the lens offset breaks the condition of holosymmetry. Instead, this system has a single plane of symmetry [36]. However, placing the system aperture at the midpoint of the transmitting and receiving lenses still provides a high degree of symmetry in the system and is therefore worth pursuing. Figure 28 depicts the off-axis interconnection setup. There is a separate aperture for

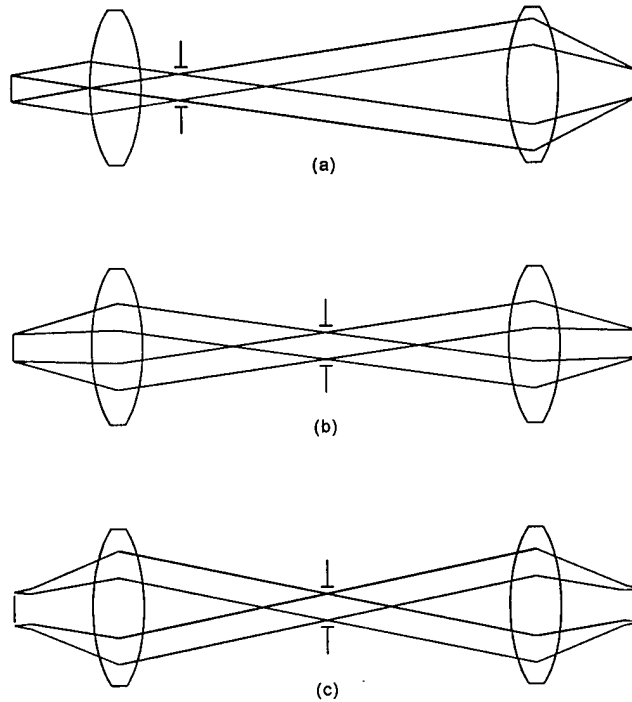


Figure 25. Depiction of VCSEL beams as they pass through the on-axis interconnection system. The VCSEL planes are on the left and the detector planes are on the right. (a) Telecentric interconnection system, (b) Symmetric interconnection system, (c) Symmetric interconnect system with auxilliary micro beam deflection elements.

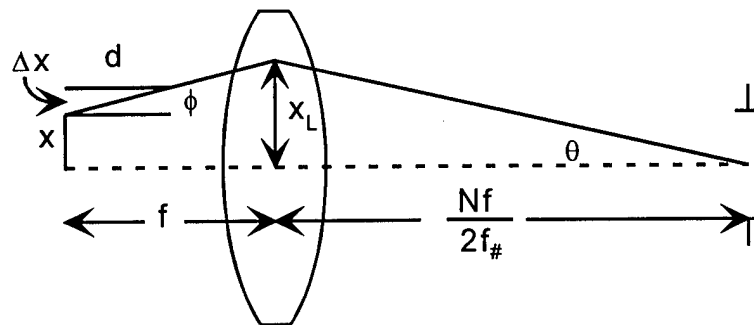


Figure 26. Geometry for deflection angle calculation.

each lens pair in the interconnection module and both clusters utilize the same region of the transmitting lens.

The geometry for analyzing the off-axis interconnection is depicted in Figure 29. The variables retain their original meanings in this figure, with the addition of: 1) the lateral distance from the lens center to the center of the cluster under examination x_c , 2) the offset from the lens

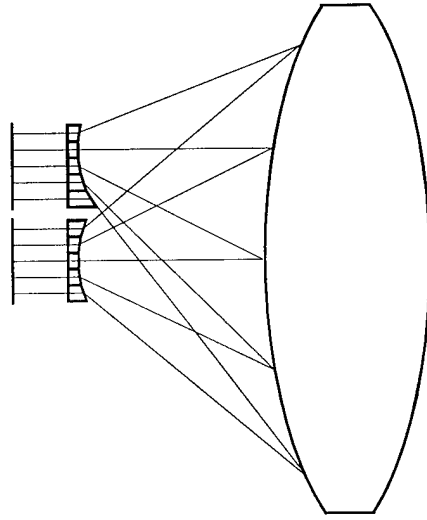


Figure 27. A collection of deflecting prisms or gratings forms a discrete negative lens.

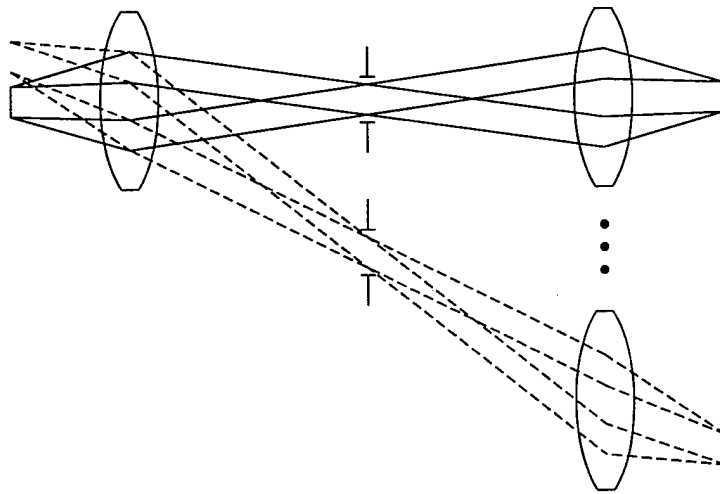


Figure 28. Multi-chip off-axis interconnection with VCSEL beam deflection to effect a central system aperture.

center to the aperture center (half the lateral distance to the receiving lens) x_{off} , and 3) there are now two beam angles θ_1 and θ_2 . The angle from the center of the cluster is θ_1 , while the angle of the beam from the element under question is θ_2 .

In this case the congruence relationships are:

$$\tan \theta_1 = \frac{x_c}{f} = \frac{2x_{off} f_{\#}}{Nf}, \text{ and} \quad (26)$$

$$\tan \theta_2 = \frac{x_c + x}{f} = \frac{2(x_{off} + x_L) f_{\#}}{Nf}, \text{ and} \quad (27)$$

$$\tan \phi = \frac{\Delta x}{d} = \frac{(x_L - x - x_c)}{f}. \quad (28)$$

Using Equations 26-28 to solve for the diffraction angle as a function of x the following is obtained:

$$\phi = \arctan \left[\frac{1}{f} \left(\frac{Nx}{2f_{\#}} - x - x_c \right) \right]. \quad (29)$$

This is the same as Equation 24, except that an angular offset proportional to x_c has been added.

Inspection of Figures 26 and 28 reveals that the effective size of the cluster is slightly increased. This effect stems from the finite distance, d , between the VCSEL and the diffraction grating. For simplicity, one can examine the on-axis case in detail. The fractional increase in cluster size is given by:

$$\frac{\Delta x}{x} = \frac{d}{f} \left(\frac{N}{2f_{\#}} - 1 \right). \quad (30)$$

Assuming $N=4$ and an $f/1$ optical system, the term in parentheses is equal to 1. The remaining term (d/f) is a small magnification – i.e., an increase on the order of 5% when $f=1\text{cm}$ and $d=0.5\text{mm}$. If the optical layout uses a regular grid pattern, this small cluster growth poses a problem. However, since the optical I/O in the proposed approach is laid out on a self-similar fractal grid geometry [14] the small magnification of cluster size does not create any overlap between adjacent clusters.

The symmetry of the new hybrid optical shuffle concept minimizes distortion – the most stringent requirement of the high-density optical interconnection module. To achieve this, the

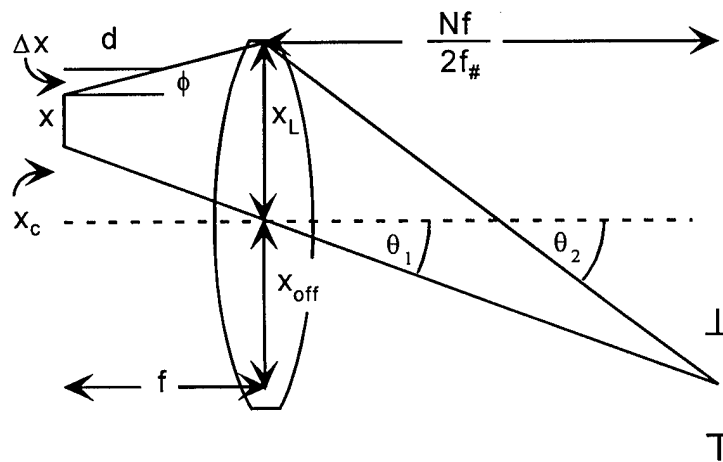


Figure 29. Geometry for off-axis analysis.

approach takes advantage of the narrow beam nature of VCSELs to effect a symmetric interconnection system for each point-to-point link in the shuffle pattern without the need for any real apertures in the system. The net result is a hybrid micro/macro approach that has optimum light efficiency and achieves high registration accuracy across the multi-chip smart pixel. The required micro-optical elements amount to a discrete negative lens above each I/O cluster. Such elements may be readily fabricated with established diffractive optical techniques. As these elements are simple gratings or micro prisms the absolute alignment of such elements is not a critical aspect of this concept. Furthermore, since resolution requirements can be easily achieved by utilizing detectors that are somewhat larger than the VCSELs (50 μm as opposed to 10 μm), the overall design of the macro-optical lenses above the array will be significantly simplified.

VI. Conclusions

This final technical report recounts the design and development of the first free-space optical interconnection based approach to handling the computational and communications complexity of high performance Viterbi Decoding and similar highly interconnected multiprocessor problems in which conventional high-speed VLSI approaches are too constraining. Multi-chip VLSI implementations are speed-limited by the power consumption, volume, and bandwidth limits of inter-chip metallic links. The overall goal of this program was, therefore, to demonstrate the feasibility of optically interconnected multi-chip parallel processing, based on the rapidly emerging *smart pixel* technology, which maintains the *on-chip speed and power efficiency* of VLSI, yet has the computational power of multiple chips.

The results of this program provide a significant step toward the incorporation of the emerging smart pixel technology into real communications-constrained applications. It is the first program to show that multi-chip smart pixel arrays can be interconnected in a high density, high bi-section bandwidth link pattern in a compact, ruggedly packaged module -- and that such modules will provide significant performance advantages. The Viterbi algorithm application has provided an important application domain -- high performance communications decoding -- for a wide range of military and commercial needs. The important advances achieved in this program have provided the basis for future efforts that will extend the reported results as the optomechanical packaging and smart pixel performance capabilities continue to grow at a rapid rate.

VII. References

- [1] M. W. Haney & M. P. Christensen, "Fundamental Geometric Advantages of Free-space Optical Interconnections," *MPPOI*, October 1996.
- [2] M. W. Haney and M. P. Christensen, "Performance Scaling Comparison for Free-space Optical and Electrical Interconnection Approaches," *Applied Optics*, Special Issue on Optics in Computing, Vol. 37, No. 14, May, 1998, pp. 2886-2894.
- [3] A. W. Lohmann, "What Classical Optics Can Do for the Digital Optical Computer," *Applied Optics*, vol. 25, pp. 1543-1549, 1986.
- [4] G. Eichmann, and Y. Li, "Compact Optical Generalized Perfect Shuffle," *Applied Optics*, Vol. 26, pp. 1167-1169, April 1987.
- [5] S.-H. Lin, T. F. Krile and J. F. Walkup, "2-D Optical Multistage Interconnection Networks," *Proc. SPIE*, vol. 752, pp.209-216, 1987.
- [6] K.-H. Brenner and A. Huang, "Optical Implementations of the Perfect Shuffle Interconnection," *Applied Optics*, vol. 27, pp. 135-137, Jan. 1988.
- [7] C. W. Stirk, R. A. Athale, and M. W. Haney, "Folded Perfect Shuffle Optical Processor," *Applied Optics*, vol. 27, pp. 202-203, 1988.
- [8] A. A. Sawchuk and I. Glaser, "Geometries for Optical Implementations of the Perfect Shuffle," *Proc. SPIE*, vol. 963, p. 270, 1988.
- [9] G. C. Marsden, P. J. Marchand, P. Harvey, and S. C. Esener, "Optical Transpose Interconnection System Architecture," *Optics Letters*, Vol. 18, No. 13, pp. 1083-1085, July 1, 1993.
- [10] R. R. Michael, M.W. Haney, & M. P. Christensen, "Experimental Evaluation of the 3-D Optical Shuffle Interconnection Module of the Sliding Banyan Architecture," September 1996, pp. 1970-1978.
- [11] H. S. Stone, "Parallel Processing with the Perfect Shuffle," *IEEE Trans. on Comp.*, C-20, 1971.
- [12] M. P. Christensen and M. W. Haney, "Two-Bounce Free-space Arbitrary Interconnection Architecture," *Proc. of the 4th International Conference on Massively Parallel Processing using Optical Interconnections*, June 1997, pp. 61-67.
- [13] M. W. Haney, M. P. Christensen, K. Raj, and P. Milojkovic, "Packaging Advantages of Macro-optical Free-space Interconnections over Micro-optical and Electrical Interconnections," *EEP-Vol. 19-1, Advances in Electronic packaging*, June, 1997, pp. 811-817.
- [14] M. W. Haney, "Self-similar Grid Patterns in Free-space Shuffle/Exchange Networks," *Optics Letters*, Vol. 18. No. 23, December 1, 1993.
- [15] M. W. Haney, "Pipelined Optoelectronic Free-Space Permutation Network," *Optics Letters*, Vol. 17, No. 4, pp. 283-285, February, 1992.
- [16] M. W. Haney and M. P. Christensen, "Sliding-banyan network performance analysis," *Applied Optics*, Vol. 36, No. 11, pp. 2334-2342, April 1997.

- [17] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*, 1992.
- [18] A. V. Krishnamoorthy, P. J. Marchand, F. E. Kiamilev, and S. C. Esener, "Grain sized considerations for optoelectronic multistage interconnection networks", *Applied Optics*, Vol. 31, No. 26, pp. 5480-5507, September 1992.
- [19] M. W. Haney and M. P. Christensen, "Optical Freespace Sliding Tandem Banyan Architecture for Self-routing Switching Networks," *Digest of the International Conference on Optical Computing*, pp. 249-250, August, 1994.
- [20] M. W. Haney and M. P. Christensen, U. S. Patent # 5,467,211, "Optoelectronic Sliding Banyan Network," issued November 14, 1995.
- [21] W.L. Hendrich, P.J. Marchand, F.B. McCormick, I. Cokgur, S.C. Esener, "Optical Transpose Interconnection System: system design and component development," Technical Digest of Optical Computing Meeting, pp 283-285, March, 1995.
- [22] C. P. Kruskal and M. Snir, "The Performance of Multistage Interconnection Networks for Multiprocessors," *IEEE Transactions on Computers* **C-32**, No. 12, pp. 1091-1098, December, 1983.
- [23] T. Nakahara, S. Matsuo, S. Fukushima, and T. Kurokawa, "Performance comparison between multiple-quantum-well modulator-based and vertical-cavity-surface-emitting laser-based smart pixels," *Applied Optics*, Vol. 35, No. 5, Feb., 1996.
- [24] A. W. Lohmann, "What Classical Optics Can Do for the Digital Optical Computer," *Applied Optics*, vol. 25, pp. 1543-1549, 1986.
- [25] C. D. Thompson, "Area-Time Complexity for VLSI," in Proceedings of the 11th Annual ACM Symposium on Theory of Computing, Atlanta, GA, April, 1979.
- [26] R. G. Schell and G. Tyras, "Irradiance form an Aperture with a Truncated-Gaussian Field Distribution," *JOSA*, Vol. 61, No. 1, January 1971, pp. 31-35.
- [27] P. Bellard and J. P. Crenn, "Changes in the characteristics of a Gaussian beam weakly diffracted by a circular aperture," *Applied Optics*, Vol. 21, No. 3, pp. 522-527, February, 1982.
- [28] L. D. Dickson, "Characteristics of a Propagating Gaussian Beam," *Applied Optics*, Vol. 9, No. 8, pp. 1854-1861, August, 1970.
- [29] F. B. McCormick, et al, "Microbeam Optical Interconnections Using Microlens Arrays," Photonic Switching, March, 1991.
- [30] M. R. Feldman, et. al., "Comparison between electrical and free space optical interconnects for fine grain processor arrays based on interconnect density capabilities," *Applied Optics*, Vol. 28, No. 18, pp. 3820-3829, September 1989.
- [31] T. Nakahara, S. Matsuo, S. Fukushima, and T. Kurokawa, *Appl. Opt.*, **35**, 860-871 (1996).
- [32] M. W. Haney, et al, *SPIE Proceedings*, Vol. 3288, January 1998, 194-203.
- [33] T. Smith, *Trans. Opt. Soc.* **23**, 139-153, 1921/1922.

- [34] G. C. Steward, *The Symmetrical Optical System*, Cambridge University Press, 1928.
- [35] R. Kingslake, *Lens Design Fundamentals*, Academic Press, 1978.
- [36] H. A. Buchdahl, *An Introduction to Hamiltonian Optics*, Cambridge University Press, 1970.